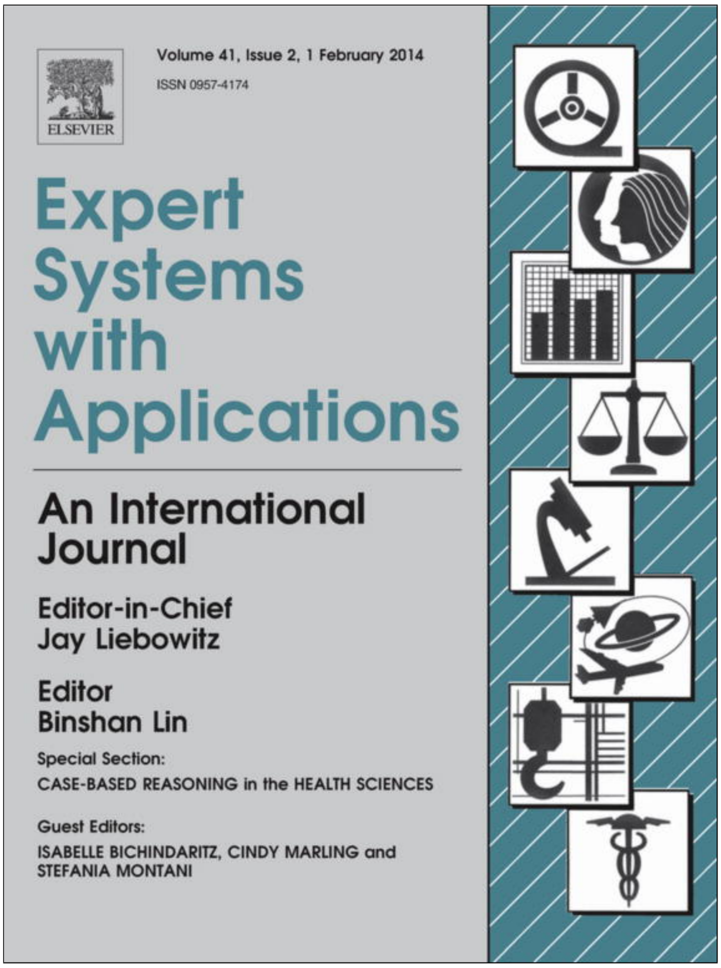
Volume 41, Issue 2, 1 February 2014
ISSN 0957-4174
ELSEVIER
Expert Systems with Applications
An International Journal
Editor-in-Chief
Jay Liebowitz
Editor
Binshan Lin
Special Section:
CASE-BASED REASONING in the HEALTH SCIENCES
Guest Editors:
ISABELLE BICHINDARITZ, CINDY MARLING and STEFANIA MONTANI

Expert Systems with Applications 41 (2014) 249–259

Contents lists available at SciVerse ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

Synergistic case-based reasoning in medical domains

Cindy Marling [a,*], Stefania Montani [b], Isabelle Bichindaritz [c], Peter Funk [d]

[a] School of Electrical Engineering and Computer Science, Ohio University, Athens, OH 45701, USA
[b] DISIT, Computer Science Institute, Università del Piemonte Orientale, 15121 Alessandria, Italy
[c] Computer Science Department, State University of New York at Oswego, Oswego, NY 13126, USA
[d] School of Innovation, Design and Engineering, Mälardalen University, SE-721 23 Västerås, Sweden

ARTICLE INFO

Keywords:
Case-based reasoning
Clinical decision support systems
Medical informatics

ABSTRACT

This paper presents four synergistic systems that exemplify the approaches and benefits of case-based reasoning in medical domains. It then explores how these systems couple Artificial Intelligence (AI) research with medical research and practice, integrate multiple AI and computing methodologies, leverage small numbers of available cases, reason with time series data, and integrate numeric data with contextual and subjective information. The following systems are presented: (1) CARE-PARTNER, which supports the long-term follow-up care of stem-cell transplantation patients; (2) the 4 Diabetes Support System, which aids in managing patients with type 1 diabetes on insulin pump therapy; (3) Retrieval of HEmodialysis in NEphrological Disorders, which supports hemodialysis treatment of patients with end stage renal disease; and (4) the Mälardalen Stress System, which aids in the diagnosis and treatment of stress-related disorders.

1. Introduction

Case-based reasoning (CBR) systems have long found fertile ground in health sciences domains (Begum, Ahmed, Funk, Xiong, & Folke, 2011; Bichindaritz & Marling, 2010; Montani, 2008). Eight international Workshops on CBR in the Health Sciences have highlighted the challenges and showcased the applications of CBR in biomedical fields. At the 2012 workshop, held at the International Conference on Case-Based Reasoning (ICCBR-12) in Lyon, several exemplary systems were featured (Lamontagne & Recio-García, 2012). In the spirit of CBR, which promotes reasoning and learning from concrete examples, four of these systems were selected as cases of medical CBR systems. These systems are:

- CARE-PARTNER, which supports the long-term follow-up care of stem-cell transplantation patients.
- The 4 Diabetes Support System, which aids in managing patients with type 1 diabetes on insulin pump therapy.
- Retrieval of HEmodialysis in NEphrological Disorders, which supports hemodialysis treatment of patients with end stage renal disease.
- The Mälardalen Stress System, which aids in the diagnosis and treatment of stress-related disorders.

In this paper, we present each of these systems in turn. Then we explore the synergies enabled and exemplified by these systems. We find a tight coupling of Artificial Intelligence (AI) research with medical research and practice. We see integration of multiple AI and computing technologies. We find that complex domains demand complex knowledge structures. We identify a need to fully leverage small numbers of available cases. We encounter time series data and develop new ways to harness it for reasoning. We integrate numeric data, including biosensor signal data, with contextual life-event data and subjective patient perceptions. In essence, the synergistic intertwining of CBR and medicine in these systems has led to new insights in both CBR research and development and medical practice. It is our hope that these experiences will be retrieved, reused, revised and retained (Aamodt & Plaza, 1994) for future CBR research and system development.

2. CARE-PARTNER

CARE-PARTNER is a decision support system for the long-term follow-up of oncology patients who have undergone stem cell transplantation (Bichindaritz, Kansu, & Sullivan, 1998). This system was built between 1996 and 2000 at the Fred Hutchinson Cancer Research Center, at the University of Washington, in Seattle. Three physicians, Keith Sullivan, Paul Martin, and Emin Kansu, and a physician assistant, Muriel Siadak, served as the domain experts. While CARE-PARTNER is no longer an active project, the ideas it pioneered have been carried over into the ongoing Mémoire project

* Corresponding author. Tel.: +1 740 593 1246.
E-mail addresses: marling@ohio.edu (C. Marling), stefania.montani@unipmn.it (S. Montani), ibichind@oswego.edu (I. Bichindaritz), peter.funk@mdh.se (P. Funk).

0957-4174/$ - see front matter
http://dx.doi.org/10.1016/j.eswa.2013.05.063

(Bichindaritz, 2006, 2007) and have influenced numerous other CBR systems in health sciences domains.

2.1. CARE-PARTNER system functionality and goals

CARE-PARTNER assists clinicians with the long-term follow-up (LTFU) of stem cell transplant patients once they have returned to their home communities. It provides online answers to questions from home care providers, who previously had to telephone nurses, who would then relay their questions to LTFU clinicians before getting back to them with clinical answers. The electronic contact management system developed to replace the old phone and paper-based system provides advantages for research and documentation and serves as an example of a medical knowledge management system.

CARE-PARTNER's decision-support is based upon proven and validated practice, helping to implement evidence-based medicine (Sullivan et al., 1997). It provides the following types of decision support:

- Interpretation for each laboratory test and procedure result.
- List of differential diagnoses, ranked by likelihood; these diagnoses are often not incompatible, since several diagnoses co-occur to cover all the signs and symptoms exhibited by the patient.
- List of steps of laboratory tests and/or procedures for diagnostic assessment.
- List of steps of planning actions for treatment.
- List of pertinent documents hyperlinked to the previous elements, such as guidelines, or textbook excerpts.

An important system requirement for CARE-PARTNER is the management of risk. A physician not specialized in a domain may not be able to critique or challenge system advice, and may not notice even severe mistakes. In the domain of stem-cell transplantation, transplant complications were quite unfamiliar to home care providers. These complications can be rapidly life-threatening, thus imposing very high standards of safety to protect patients. Therefore, the reliability and safety of the system were of paramount importance.

2.2. CARE-PARTNER system design

Fig. 1 shows CARE-PARTNER's reasoning cycle. This multimodal reasoning cycle combines case-based reasoning, rule-based reasoning, and information retrieval. CARE-PARTNER's reasoning steps are generalizations of the steps defined in these respective methodologies. The cooperation of the different knowledge sources is driven by the LTFU domain, in which, as in most medical domains, knowledge takes several forms:

1. **Practice guidelines:** A practice guideline is composed of systematically developed textual statements, designed for practitioners and patients, which will be helpful in making clinical decisions on the prevention, diagnosis, treatment and management of selected conditions. Guidelines are represented as rules that are embedded in prototypical cases.
2. **Practice pathways:** A practice pathway covers the same type of knowledge elements as a practice guideline, but it is specialized to the LTFU domain. While practice guidelines are represented via text, practice pathways are expressed in the knowledge representation formalism of the decision support system. Practice pathways were created by a group of LTFU experts exclusively for the CARE-PARTNER system. Pathways correspond to prototypical cases, and they are represented as cases in the system.
3. **Practice cases:** A practice case is an example of a problem-solving situation as solved by an expert or possibly a group of experts. It is essentially a real patient problem-solving situation, and not a prototypical one as for a practice guideline or a practice pathway. It is represented as a case in the system.
4. **Medical textbooks:** A collection of documents serves as documentation and explanation during the reasoning process, often in hyperlinked form.

Intensive knowledge elicitation efforts were required to build the case base around a knowledge base of the domain. It was determined early on in the project that cases were not available in electronic format at a level of detail required for CBR. For instance, the patient database did not include patient treatments, or most of the signs and symptoms, but only the main events abstracted from the paper charts. The project team had to come up with prototypical cases to bootstrap the system, which took over two years to develop at a level of thoroughness and consistency needed to achieve high quality decision support.

This system was unique because its proposed recommendations spanned not only diagnosis, but also lab result interpretation, and treatment planning. An extensive ontology was developed including 1109 diseases, 452 functions (also known as signs and symptoms), 1152 labs, 547 procedures, 2684 medications, and 460 sites. Most of the terms naming these objects were standardized using the Unified Medical Language System (UMLS) semantic network (National Library of Medicine, 1995). Only terms not corresponding to objects in the UMLS were added to the domain specific ontology. In particular, the planning actions used in the treatment part of a prototypical case did not exist in the UMLS and were all created for the system.

The cornerstone of the knowledge acquisition process was the conception of prototypical cases, or clinical pathways. The 91 implemented clinical pathways primarily correspond to clinical diagnostic categories, with some of them corresponding to essential signs and symptoms requiring specific assessment or treatment actions. The clinical pathways are knowledge structures represented using the ontology described above. A prototypical case comprises three parts:

1. A list of findings, corresponding to signs and symptoms.
2. A diagnosis assessment plan, which is a plan to follow for confirming (or informing) the suspected diagnosis.
3. A treatment/solution plan, which is a plan to follow for treating this disease when confirmed, or a solution when the pathway does not correspond to a disease.

The diagnosis assessment part and the treatment part of the case can also be viewed as simplified algorithms, since they use *if-then-else* structures, *loop* structures, and *sequence* structures of actions in time. When instantiated with an actual patient's data, this provides a diagnosis assessment plan or treatment plan tailored to the specific patient. In this way, the knowledge structure allows for sophisticated adaptation when reusing a prototypical case.

2.3. CARE-PARTNER evaluation

Table 1 shows the results of an evaluation in which two expert clinicians rated the system using the scale *Fails to Meet Standards/Adequate/Meets All Standards* (Bichindaritz, 2006). This evaluation covered 163 different clinical situations or cases, corresponding to contacts between the system and a clinician, for three patients. As seen in Table 1, the system was rated 82.2% of the time as *Meets all Standards* and 12.3% of the time as *Adequate*, for a total of 94.5%

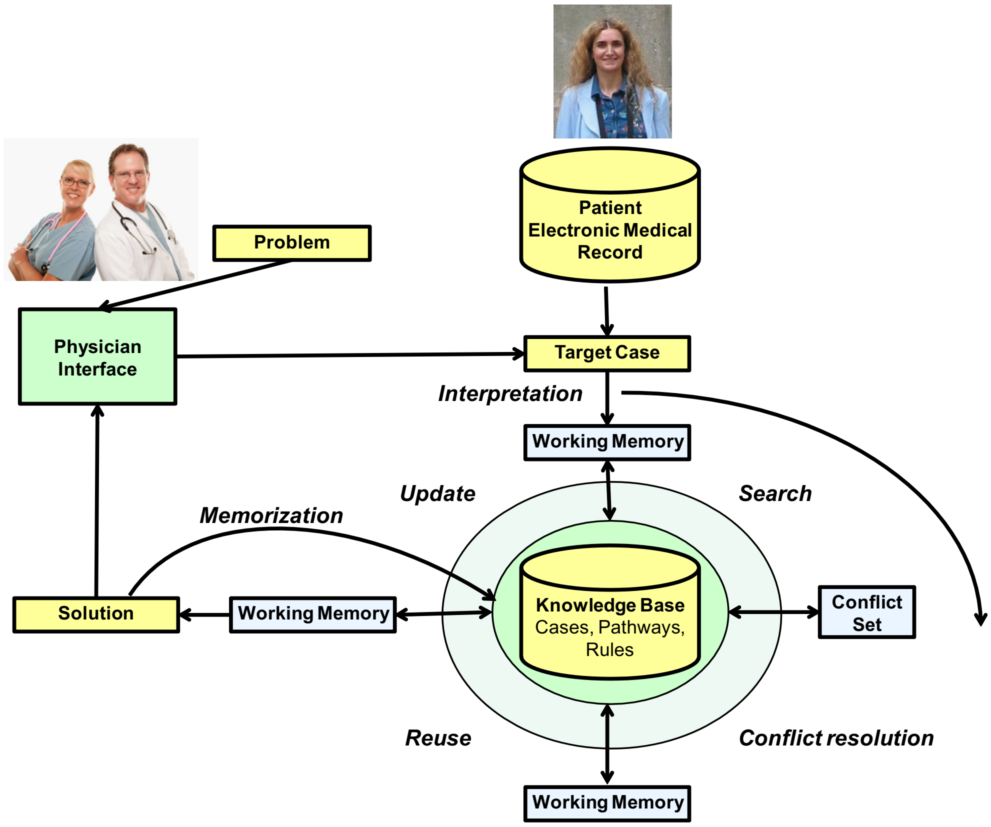

Fig. 1. CARE-PARTNER's Reasoning Cycle.

Table 1
CARE-PARTNER system evaluation results.

Type of case	Number of cases	Fails to meet standards (%)	Adequate (%)	Meets all standards (%)
Labs	54	3.7	3.7	92.6
Procedures	67	8.9	3.0	88.1
Diagnosis	68	16.2	13.2	70.6
Treatment	71	9.9	11.3	78.8
Pathways	47	8.5	8.5	83.0
Total	163	5.5	12.3	82.2

of results judged to be at least clinically acceptable by the medical experts. Table 1 also shows that the advice provided by the system covers most of the clinicians' tasks: lab results interpretation, procedure results interpretation, diagnosis, treatment, and pathways retrieval.

Another part of the evaluation dealt with measuring the progress of the system when solving new contact cases. This ability of the system to learn was evaluated on the complete charts of three different patients. The performance of the system significantly improved between patients 1 and 3 to reach 98.6% acceptability for the 54 contacts in the third patient's chart (Bichindaritz, 2006). Since the Therac-25 accidents had recently shown that not even 100% reliability was sufficient to ensure patient safety (Leveson & Turner, 1993), the system was further extended to include a safety control module capable of referring cases requiring particular attention to direct clinician supervision.

3. The 4 Diabetes Support System (4DSS)

The 4 Diabetes Support System (4DSS) aims to assist patients with type 1 diabetes (T1D) and their professional caregivers (Marling, Shubrook, & Schwartz, 2009, 2011a, 2012). Work on this system began in 2004 and continues to this day. There are two domain experts: Frank Schwartz, an endocrinologist; and Jay Shubrook, a diabetologist. They are practicing clinicians, who have treated several hundred T1D patients, as well as faculty members of the Ohio University Heritage College of Osteopathic Medicine.

3.1. The diabetes management domain

There are an estimated 346 million people worldwide who have diabetes (World Health Organization, 2012). Approximately 20 million of them have type 1 diabetes, the most severe form, in which the pancreas fails to produce insulin. Because insulin is an essential hormone needed to convert food into energy, T1D patients depend upon external supplies of insulin. The T1D patients involved in the 4DSS project are treated with insulin pump therapy. An insulin pump continuously infuses the patient with basal insulin. The patient may instruct the pump to deliver additional insulin boluses to account for meals or blood glucose excursions.

While diabetes can not yet be cured, it is actively managed through blood glucose control. Good blood glucose control is known to help delay or prevent long-term diabetic complications, including blindness, amputations, kidney failure, strokes, and heart attacks (Diabetes Control & Complications Trial Research Group, 1993). Effective blood glucose control entails vigilant self-monitoring of blood glucose levels. T1D patients prick their fingers from 4 to 6 times a day and use glucometers to measure their blood. They may also wear continuous glucose monitoring (CGM) devices, which produce blood glucose readings every 5 min.

Blood glucose monitoring data is relayed to physicians, who must manually interpret it to find blood glucose control problems and recommend appropriate therapeutic adjustments. While voluminous blood glucose data contributes to data overload for physicians, data concerning life events that impact blood glucose levels is not routinely maintained. Physicians may feel, paradoxically, that they have too much data and yet not enough data at the same time.

3.2. 4DSS system overview

The 4DSS is a hybrid case-based reasoning (CBR) system that detects problems in blood glucose control and suggests personalized therapeutic adjustments to correct them. Fig. 2 shows a graphical overview of the 4DSS. The system is data driven. Blood glucose data comes from glucometers and CGM devices. Insulin data comes from the patient's pump. The patient uploads blood glucose and insulin data to Medtronic's proprietary CareLink system (Medtronic, 2012), where it is extracted and transferred to the 4DSS database. The patient manually enters data about life events that impact blood glucose levels, including food, exercise, sleep, work, stress and illness. Originally entered via computer-based browsers, life-event data is now entered via smart phones.

The situation assessment module scans patient data. Traditionally in CBR systems, situation assessment begins with a static description of the current problem. It is different in this domain, because blood glucose control problems continue over time, and because patients are not necessarily aware of problems when they occur. The 4DSS situation assessment module has three major components: problem detection, glycemic variability classification, and blood glucose prediction. These components were built using rule-based reasoning, machine learning algorithms, and time series prediction techniques, giving the 4DSS its hybrid character.

The problem detection component contains 18 rule-based routines that incorporate physician strategies for finding problems in patient data. At a high level, these routines look for problems involving: (1) hyperglycemia, or high blood glucose, which contributes to long-term diabetic complications; (2) hypoglycemia, or low blood glucose, which may result in severe immediate reactions, including weakness, dizziness, seizure or coma; (3) fluctuations between hyper-and hypoglycemia; and (4) lapses in diabetes self-care.

The glycemic variability classication component assesses problems involving blood glucose fluctuation. It detects the problem of excessive glycemic variability, which is believed to presage long-term complications caused by oxidative stress. When expert rules proved inadequate for detecting this problem, machine learning algorithms, including multi-layer perceptrons and support vector machines, were introduced. These algorithms classify 24-h blood glucose plots by variability level to match physician gestalt perception of such plots.

The blood glucose prediction component, currently under construction, incorporates time series prediction techniques to anticipate problems before they occur. While blood glucose data is not currently available in real time and must be scanned retrospectively, we are preparing for its near-term availability. Predicting problems even 30 min in advance would give patients time to take preventative actions, enhancing patient safety.

The situation assessment module reports the problems it finds to a physician, who must then select problems of interest. A selected problem triggers the case retrieval module of the 4DSS. The case retrieval module obtains the most similar cases from the 4DSS case base.

The case base includes 80 cases, each of which contains a specific blood glucose control problem experienced by a T1D patient, a physician's recommended therapeutic adjustment for that problem, and the clinical outcome for the patient after making the therapeutic adjustment. These cases were compiled during clinical research studies in which: (1) T1D patients contributed blood glucose, insulin and life event data; (2) physicians manually identified

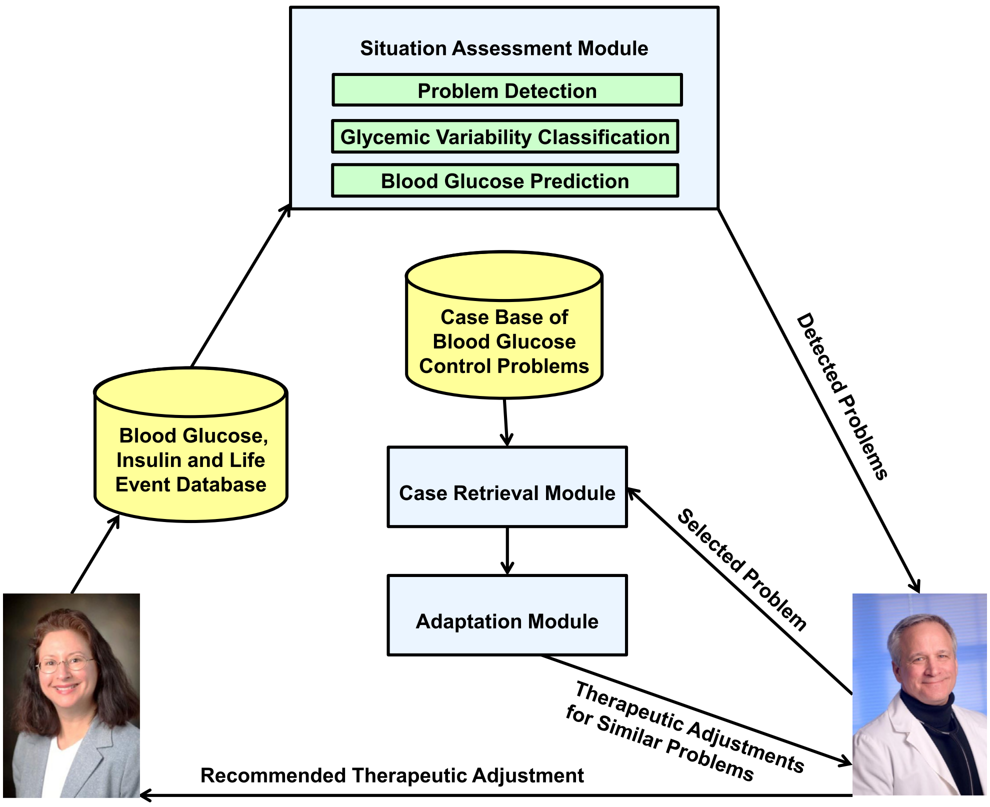

Fig. 2. Overview of the 4 Diabetes Support System.

blood glucose control problems and recommended solutions to patients; (3) patients made the recommended therapeutic adjustments (or not); and (4) physicians examined subsequent patient data to determine the efficacy of the solutions.

To retrieve the most similar cases from the case base, the case retrieval module employs a traditional two-step process. First, a subset of potentially similar cases is identified, and then the nearest neighbors are selected from that subset. In the first step, cases are partitioned by problem type. In the second step, a standard nearest neighbor metric is used. Domain specific similarity functions are combined with empirically determined weights to obtain an overall score for each case. Cases scoring above a similarity threshhold are forwarded to the adaptation module.

The adaptation module personalizes a solution from a retrieved case to fit the situation of a current patient. It begins with the most similar case, but if the solution in that case is not adaptable, it considers the next most similar case, and so on. A solution is a therapeutic adjustment composed of one or more actions that a patient can take. During adaptation, individual actions may be deleted or modified. For example, one possible action is to have a bedtime snack. If the current patient is already having an adequate bedtime snack, this action could be removed from the recommendation. In other situations, the advice could be modified so that the patient eats more or less food before bed, eats a different type of food before bed, or has a snack at a different time of day.

The adapted therapeutic adjustment is relayed to the physician as decision support. The physician decides whether or not to relay the recommendation to the patient. It has long been a goal to provide low-risk advice directly to patients, in real-time, as well as to their physicians. However, this must remain a future goal until the safety and efficacy of the system is proven through clinical trials and approval is obtained from governmental regulatory agencies.

3.3. 4DSS evaluation

Each component of the 4DSS is evaluated upon completion. These evaluations have provided proof of concept, illuminated system strengths and weaknesses, and guided system development. Note that a definitive clinical trial, assessing system impact on patient outcomes, remains to be conducted.

The problem detection component was evaluated after the first and second 4DSS clinical studies. In the first evaluation, a panel of diabetes practitioners rated a sampling of problem detections, and in the second, each patient's physician rated all of the problem detections for the patient. In the first test, 77.5% of problem detections were rated as correct (Marling, Shubrook, & Schwartz, 2008), while in the second, 97.9% were rated as correct (Schwartz, Vernier, Shubrook, & Marling, 2010).

The glycemic variability classification component was also evaluated twice. Here, ten-fold cross validation was used to determine the accuracy, sensitivity and specificity of each potential classifier, where correctness is defined as matching physician classifications. In an early test, a naive Bayes classifier matched physicians 85% of the time (Marling, Shubrook, Vernier, Wiley, & Schwartz, 2011b). The current best classifier, a multi-layer perceptron, has accuracy, sensitivity and specificity of 93.8%, 86.6%, and 96.6%, respectively (Wiley, Bunescu, Marling, Shubrook, & Schwartz, 2011).

The case retrieval module was evaluated by a panel of diabetes practitioners after the first and second clinical studies. Leave-one-out testing was used to provide a sampling of case retrievals for evaluation. In the first test, evaluators rated the retrieved cases as similar to test cases 80% of the time and rated the retrieved solutions as beneficial for test patients 70% of the time (Marling et al., 2008). In the second test, they rated retrieved cases as similar 79% of the time and retrieved solutions as beneficial 82% of the time (Marling et al., 2011b).

The adaptation module was evaluated by showing physicians sample problems, with both original and adapted solutions, and eliciting feedback on a questionnaire. Physicians rated the original solutions as being fine without adjustment 47% of the time, needing minor adjustment 40% of the time, and needing major adjustment 13% of the time. They judged the adapted solutions to be better than the original solutions 83% of the time.

4. Retrieval of HEmodialysis in NEphrological disorders (RHENE)

Retrieval of HEmodialysis in NEphrological Disorders (RHENE) supports physicians working in the domain of end stage renal disease (ESRD). Work on RHENE began in 2004 and continues to this day. The expert for RHENE is Roberto Bellazzi, a nephrologist at the Vigevano Hospital in Italy.

4.1. The end stage renal disease domain

ESRD is a severe chronic condition that corresponds to the final stage of kidney failure. Without medical intervention, ESRD leads to death. Hemodialysis is used to treat ESRD patients. During hemodialysis, an electromechanical device called a hemodialyzer clears the patient's blood of metabolites, re-establishes acid–base equilibrium and removes excess water. A single hemodialysis session typically lasts for four hours. On average, a hemodialysis patient receives three treatment sessions per week. The hemodialyzer tracks several time series variables during each session, sampling each at intervals of from 1 to 15 min. These variables are analyzed to assess the efficacy of the hemodialysis treatment session and to ensure that the patient's treatment adheres to his or her therapy plan.

Interpreting a hemodialysis session as a *case*, we must deal with cases with time series features. Interpreting time series features on screen or on paper can be tedious and prone to errors. Physicians may be asked to recognize small or rare irregularities in the series itself, or to identify partial similarities with past patient situations, which do not depend upon the individual values in the series. While extremely important for patient care, such identification crequires a significant amount of expertise. Therefore, having an automated data interpretation and decision support system is desirable.

4.2. RHENE system overview

In time dependent domains, the need to describe process dynamics strongly impacts both case representation and case retrieval, as analyzed in (Montani & Portinale, 2006). Most reported approaches to similarity-based time series retrieval are founded on the common premise of dimensionality reduction, which simplifies knowledge representation (see the survey in (Hetland, 2003)). Dimensionality is often reduced by means of a mathematical transform – such as the Discrete Fourier Transform (Agrawal, Faloutsos, & Swami, 1993) – able to preserve or underestimate the distance between two time series. However, mathematical transforms have several limitations, as they can be computationally complex, and usually work in a black box fashion with respect to end users. In contrast, RHENE (Montani, Leonardi, Bottrighi, Portinale, & Terenziani, 2011) implements a framework for time series retrieval that exploits Temporal Abstractions (TA) (Shahar, 1997) to reduce time series dimensionality, with multi-dimensional index structures to make retrieval efficient.

Using TA allows greater interpretability of the output results and understandability of the retrieval process. It maps huge amounts of temporal information to a compact representation, by aggregating adjacent time series points sharing commonalities (e.g., the same qualitative level, the same trend direction) into a single interval, labeled by a symbol. This technique not only summarizes the original longitudinal data, but also highlights meaningful data characteristics in a clear, symbolic, high level view. The exploitation of TA for case exploration and retrieval, as well as for data interpretation, is, to the best of our knowledge, a novel contribution of RHENE. TA are typically applied for data interpretation, but not for case retrieval.

The basic principle of TA methods is to move from a point-based to an interval-based representation of the data (Bellazzi, Larizza, & Riva, 1998), where the input points (events) are the elements of the time series, and the output intervals (episodes) aggregate adjacent events sharing a common behavior, persistent over time. Episodes are identified by symbols. Basic abstractions can be further subdivided into state TA and trend TA. State TA are used to extract episodes associated with qualitative levels of the monitored feature, e.g., low, normal, or high values. Trend TA are exploited to detect specific patterns, such as increase, decrease or stationary behavior, from the time series.

RHENE supports multi-level abstractions of the original data. Time series values can be abstracted and queried at finer or coarser levels of detail, along two dimensions: a taxonomy of symbols, and a taxonomy of time granularities. For instance, a taxonomy of trend symbols can be introduced, in which the symbol *I* (increase) is further specialized into I_W (weak increase) and I_S (strong increase), according to the slope. As another example, a series of two adjacent intervals of I_W, each having a duration of half an hour, can be merged into a single I_W interval, with a duration of 1 h. To scale up from two or more values expressed at a finer granularity to a single value expressed at a coarser one, an *up* function is provided. This function is domain dependent, but obeys some general constraints, including "persistence" (the result of coarsening two granules with the same symbol *x* is a larger granule still labeled as *x*) and "monotonicity" (ordering among symbols, if any, is preserved) (Montani et al., 2011).

Once abstracted, two time series can be compared to each other, enabling the retrieval of similar cases. Retrieval depends upon a distance metric that measures the similarity between symbols in the taxonomy. Different distance functions can be employed, as long as the distance of each symbol from itself is zero and other distances are consistent with respect to any symbol ordering.

A query language was developed to facilitate case retrieval. Queries are expressed as sequences of symbols at different levels of detail. RHENE supports ground queries, i.e., queries composed of symbols at the lowest abstraction level in both taxonomies, and abstract queries, including those with symbols at different levels in the symbol taxonomies. Queries as regular expressions are also supported, making the query process flexible and user friendly.

To increase retrieval efficiency, RHENE includes an indexing strategy, which exploits bi-dimensional taxonomic indexes. A forest of index structures provides a flexible indexing of cases at different levels of the symbol and/or time granularity taxonomies. The root node of each index structure is represented by a string of symbols, defined at the highest level in both dimensions. An example is shown in the right box of Fig. 3. Here, the root node *I* is refined along the leading time dimension from the 4-h to the 2-h granularity, so that the nodes *I I*, *I S* and *S I* stem from it, provided that $up(I,S) = I$ and $up(S,I) = I$, where *S* stands for stationarity. From each node of the leading dimension structure, another index can stem, keeping the time granularity abstraction level fixed. The index develops orthogonally with respect to the leading dimension.

In summary, RHENE operates in the following manner, as illustrated by Fig. 3:

- Time series features in ESRD cases (i.e., hemodialysis sessions) are pre-processed by a TA server at the ground level of the symbol and time granularity taxonomies.
- Pre-processed cases are then stored in a relational database.
- When the physician needs to evaluate a patient, he or she can query the database to retrieve similar cases at any level of detail.
- Indexes enable quick and interactive query answering.
- Retrieved cases are then shown to the physician as decision support. The physician remains responsible for making decisions regarding the patient's therapy.

4.3. RHENE evaluation

RHENE was experimentally evaluated, using a dataset of 10,388 actual hemodialysis sessions (i.e., ESRD cases) collected at the Vigevano Hospital in Italy. The TA approach was compared to the more classical approach implemented in an earlier version of RHENE (Montani, Portinale, Leonardi, Bellazzi, & Bellazzi, 2006). That approach was based on DFT for dimensionality reduction and on spatial indexing through TV-trees for improving retrieval performance. Quantitative experiments were run to compare the query answering time required by the two approaches, as well as their scalability when dealing with a case base progressively growing in size. Subsets of from 2,000 to 10,388 actual cases were used to evaluate each system's performance. As shown in Fig. 4, the TA-based method proved to be much more efficient in query answering than the DFT-based method, with query times under 1 s.

A qualitative comparison of the two approaches was also made, through examination of individual case studies. Details of this examination are available in (Montani et al., 2011). In brief, it was found that the ability to consider trends was advantageous for TA-based retrieval. DFT sometimes missed relevant cases, because it considers only point-to-point distances between cases, looking for the best overall alignment. Shifts along an axis, such as lower absolute values, could prevent DFT from finding cases with similar trends. Furthermore, the TA-based approach provided the user with more easily understood queries and results than DFT's more mathematical black box approach.

5. The Mälardalen Stress System (MSS)

The Mälardalen Stress System (MSS) provides decision support for the diagnosis and treatment of stress (Ahmed, Begum, & Funk, 2012, 2011). This system was built between 2002 and 2011 at Mälardalen University in Västerås, Sweden. The experts for this system were psychologists Bo von Schéele and Erik Olsson. Dr. von Schéele has over 30 years' experience in clinical stress diagnosis and has pioneered new biosensor and biofeedback methods.

5.1. The stress management domain

Stress is a factor of daily life that can adversely impact health and wellbeing. While the causes of stress can not typically be eliminated, patients can be taught to effectively deal with stress, minimizing its ill effects. A clinical measurement of stress is Finger Temperature (FT). By placing a sensor on a patient's fingertip, a continous temperature profile can be obtained. In general, lower finger temperatures are indicative of greater stress. However, changes in FT vary from patient to patient, and correctly interpreting and analyzing them requires knowledge and experience. Moreover, additional factors, such as the patient's feelings, behaviors, environment and lifestyle, also play a role in stress management. Contextual features, indicating a patient's perception of stress, are collected as text and also via Visual Analog Scale (VAS) input.

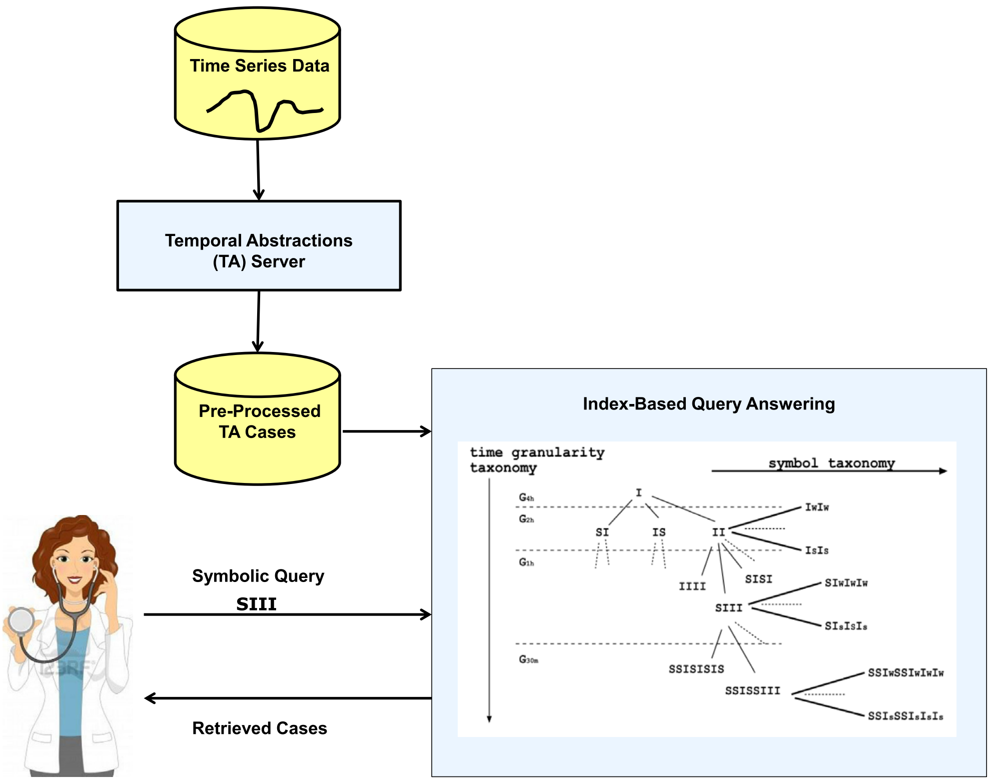

Fig. 3. RHENE System Overview.

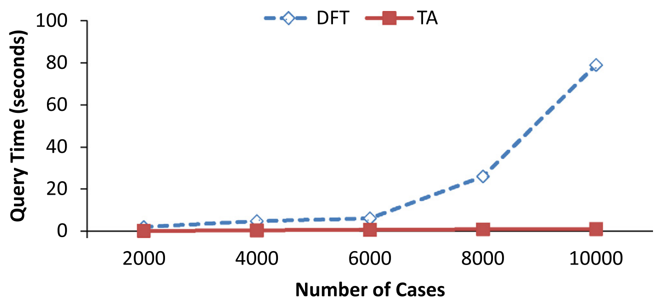

Fig. 4. Average time, in seconds, to answer a query as the number of cases increases, for DFT and TA.

VAS is used to measure subjective characteristics or attitudes on a scale of 0 to 10. Both sensor signals and textual information are used to diagnose a patient's level of stress.

Biofeedback training is used as treatment to control stress. During biofeedback training, a patient can alter his or her physiological or psychological state while observing the changes in FT on a graph. As the patient sees how psychophysiological change is represented on the graph, he or she can train the body and/or mind to change the biological response to stress. Relaxation exercises may be recommended to aid the patient in making positive changes.

5.2. MSS system overview

The MSS demonstrates how some common CBR techniques need to be modified in order to provide clinicians with effective decision support based on past cases relevant to the current patient. This clinical usage adds a number of specific requirements. For example, in addition to considering case similarity, it is important to consider case outcomes. A case with a severe outcome may be very relevant, as a clinician may need to take precautions to avoid a similar result. An example is that most patients' finger temperatures decrease during stress and increase during relaxation, a normal reaction from the parasympathetic nervous system. But in some patients, the effect on finger temperature may differ due to rare conditions the clinician must consider during diagnosis. If the current patient is similar to many past patients diagnosed as not at risk for stress-related illnesses and similar to one patient having a condition with severe health consequences if left untreated, the latter case is of most interest to the clinician, who may take additional measurements to rule this case out.

To meet the needs of the medical domain, the case has been structured with the following components: (a) symptom features, including sensor readings and contextual data; (b) diagnosis; (c) action taken, i.e., treatment and biofeedback training; and (d) outcome, e.g., how quickly the patient improves. The case may also contain (e) comments, which can be added by a clinician who wishes to share perceptions and/or important references with colleagues. Comments may be added by the clinician responsible for the case or by a domain expert.

The symptom features for a case are derived from a calibration protocol in which the patient is asked to complete a series of stressful, relaxing and neutral tasks while FT measurements are collected (Ahmed, Begum, Funk, Xiong, & von Schéele, 2011). This calibration phase is necessary, because sensor readings vary widely from patient to patient. Sensor readings that are normal for one person may be alarming for another individual. The patient is also asked to supply contextual data via text and VAS input. A total of 25 sensor and contextual features comprise this section of a case. The diagnosis part of the case holds the stress classification of the patient, which may be: VeryRelaxed, Relaxed, Normal/Stable, Stressed or VeryStressed. A confidence level for the classification, denoted as High, Medium or Low, is also recorded for each diagnosis. The actions taken comprise the treatment, and the outcome is the result of the treatment, i.e., the degree of improvement after treatment.

Biofeedback is a tool used by clinicians in treatment and overall stress reduction. It is included in the section of the case for action taken. The latest diagnosis is used as one of the parameter settings for biofeedback. The patient gets feedback on parameter improvement. Another important parameter is recovery time after stress, i.e., the time it takes for a patient to recover after a stress-phase. This parameter can be used to select relaxation exercises/treatment for the patient. During the exercise, the biofeedback system monitors the patient and calculates the recovery time after stress. This enables patients to practice biofeedback either with a clinician at hand or alone with a computer system and sensors. The relaxation exercise/treatment is selected to achieve good results, based on how good the effect was on previous similar patients. If the recovery time does not change with the selected exercise, then a different exercise is selected. Exercise selection is based on past cases and the result of the exercise/treatment. Different exercises have different effects on patients. Selecting exercises and following patient progress in this way increases the efficiency of the exercise.

The MSS supports stress diagnosis and treatment in three phases: (1) analyze and classify a patient and make a risk assessment; (2) determine individual levels and parameters; and (3) adapt and conduct biofeedback training. CBR is used as the core technology for the system, as illustrated in Fig. 5. Knowledge discovery contributes part of the domain knowledge used in case retrieval. Since patients in the case library have already been diagnosed/classified, these cases can be used to identify the features that are most important for case comparison and retrieval (Funk & Xiong, 2006). Other AI techniques, including fuzzy logic, rule-based reasoning and textual information retrieval, are also incorporated in supporting roles.

5.3. MSS evaluation

The MSS system has been evaluated using a number of different methods. Diagnostic performance has been compared with that of clinicians, including clinical experts and junior clinicians with limited practice in diagnosing stress using sensor readings. Using multiple test sets, the system was able to classify over 80% of cases correctly, compared to between 57% and 69% for junior clinicians and 73% for a senior clinician (Ahmed et al., 2012). In comparing system outputs to clinician diagnoses, it was important to consider the consistency of clinician responses when confronted with the same patient more than once. In another test of diagnostic performance, leave one out cross validation was used. In this evaluation, cases were removed from the case library one at a time to see how well the system could classify each one. Using a fuzzy matching algorithm, the MSS was able to perform close to expert level, as reported in (Begum, Ahmed, Funk, Xiong, & von Schéele, 2009). In addition to judging overall accuracy, it was important to consider the sensitivity and specificity of the system. From a clinical perspective, missing a stressed patient (false negative) is less acceptable than identifying a healthy individual as stressed (false positive). A false positive would likely be identified as such before beginning treatment, based on additional pre-treatment evaluation.

6. Research synergies and trends

In this section, we elaborate on the research synergies and trends highlighted by CARE-PARTNER, the 4DSS, RHENE, and the MSS.

6.1. Coupling AI research with medical research and practice

All four systems were built through strong collaborations with medical researchers and clinicians. While we have focused above on AI research and development, there has been a tight coupling with medical research and practice, as well. For example, CARE-PARTNER supported *evidence-based practice* for oncology. At the time of the original study, stem cell transplantation was a very new procedure, and there were no established guidelines for the long-term follow-up care of patients after they left the cancer center. By compiling cases of patient problems, solutions, and outcomes, CARE-PARTNER aided in defining best clinical practices. The 4DSS, RHENE, and the MSS can all be characterized as promoting the practice of *personalized medicine*. These systems look at how individual patients respond to diabetes therapy, hemodialysis,

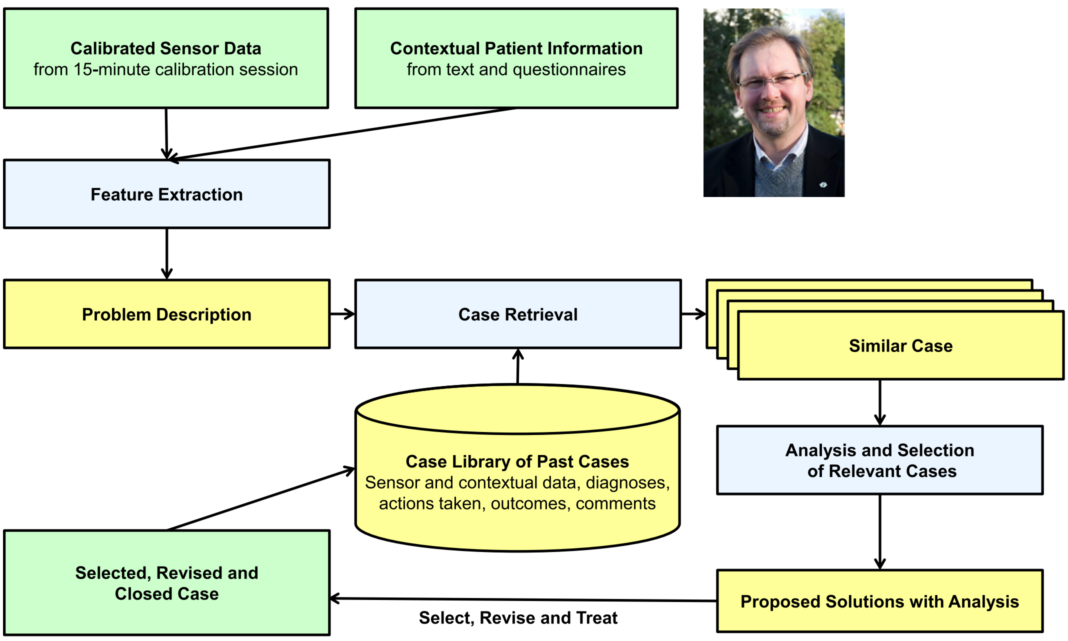

Fig. 5. Overview of the Mälardalen Stress System.

and stress, respectively, and suggest personalized therapeutic adjustments based on individual patient needs.

The 4DSS contributes to medical research in the areas of glycemic variability (GV) measurement and blood glucose prediction. Because there was no accepted metric for GV in routine clinical use, but 4DSS experts wanted to detect excessive GV, new metrics were developed for the 4DSS. These metrics were subsequently published and presented to the diabetes technology community (Schwartz et al., 2010), and the 4DSS project is now leading the way in developing a consensus GV metric for routine clinical use. Current research on blood glucose prediction not only aids in intelligent decision support, but potentially aids in the development of an "artificial pancreas," a project of the Juvenile Diabetes Research Foundation.

RHENE supports medical practice and research as well. Hemodialysis sessions are ordinarily judged only on the basis of macroscopic observations of time series features. RHENE provides a deeper insight into the clinical situation, highlighting types of anomalies which, if not leading to immediate hemodialysis failure, could produce poor therapeutic results in the long run. (See, for example, the experiments in Montani et al. (2006)). Moreover, the use of RHENE may lead to the identification of systematic mappings between TA trend behaviors (e.g., an increase in diastolic pressure) and specific pathologies or complications, in an application domain where this kind of knowledge does not currently exist.

Expert clinicians with the MSS project found that their work on the system helped them to refine their own methods and approaches for diagnosing patients. They also found that the systematic analysis of symptoms required by the system improved the quality and value of the electronic health records (EHRs) that they maintained for their patients. The final outcomes of treatments had not always been recorded in detail prior to system development. As tracking outcomes was essential for case-based reasoning, the EHRs for patients of participating physicians became more complete. Furthermore, the close collaboration made it more explicit how clinical experts, while reading symptoms and measurements, would recognize similarities to past patients that would aid in diagnosis and treatment. This led to the discovery of new features that were valuable in the diagnosis process (Funk & Xiong, 2006, 2008).

6.2. Integrating multiple AI and computing methodologies

All four systems are multi-modal, or hybrid, systems, which synergistically combine CBR with other AI and computing methodologies. CARE-PARTNER combines CBR with rule-based reasoning and information retrieval. Its knowledge base encompasses both theoretical knowledge and experiential knowledge, which may be expressed either in a controlled vocabulary or textually. The 4DSS integrates rule-based reasoning for problem detection and case adaptation, multiple machine learning algorithms for glycemic variability classification, and support vector regression for blood glucose prediction. RHENE incorporates temporal abstractions for representing, comparing, and retrieving cases. The MSS leverages rule-based reasoning, fuzzy logic, and information retrieval. Fuzzy rule-based classification is used when the domain knowledge is precise (e.g., if X and Y are high, then we know that Z is high) but the terminology used is not precise (e.g., the same absolute value may be high or low, depending on context). Information retrieval is used to handle the free text given in the case, e.g., the patient history.

Driven by the demands of specific medical domains, these integrations push the envelope of how knowledge-based systems can be engineered. For example, hemodialyzers generate data that could not be readily reasoned about without some form of feature reduction or abstraction. So, RHENE introduced a novel combination of TA and CBR that could be extended to other domains having time series data. Similarly, CARE-PARTNER's early information retrieval integration has become increasingly relevant as information from electronic health records and large medical corpora become increasingly available.

6.3. Leveraging small numbers of available cases

In this era of Big Data, there is much emphasis on extracting useful information from large volumes of available data. For medical decision support applications, however, it is still unusual to have all of the relevant data at hand. Although RHENE did have thousands of cases to work with, CARE-PARTNER, the 4DSS, and the MSS effectively leveraged small numbers of available cases. CARE-PARTNER, for example, began with an incomplete patient database containing extracts from paper-based patient charts. Ninety-one prototypical cases were built through extensive knowledge engineering efforts, enabling domain knowledge to supplement the available data. The 4DSS began with an abundance of blood glucose data, but the contextual life events needed to interpret it were not routinely maintained. A series of three clinical research studies was conducted, in which 80 problem/solution/outcome cases were developed from actual patient data. Each case contains over 140 data elements, hierarchically organized, to represent a problem solving experience. The MSS project conducted clinical studies to capture and represent patient stress profiles. However, case coverage was initially incomplete, because some profiles, such as the Very Relaxed profile, were not well represented. Artificial cases were generated to augment the actual patient cases by using generalized sensor features and a fuzzy inference system. This improved the system's ability to retrieve applicable cases (Ahmed, Begum, Funk, & Xiong, 2009). While there are many ways to leverage large numbers of exemplars, the ability to solve problems with just a few knowledge-rich cases is a strength of the CBR approach.

6.4. Reasoning with time series data

All four systems reason about cases that develop over time. This is in contrast to most CBR systems, in which a snapshot of current conditions is sufficient for assessing and reasoning about the problem at hand. CARE-PARTNER, for example, is for *long-term* follow-up care, and the patient's past history informs current care.

Time series data for the 4DSS comes in the form of blood glucose sensor data, collected for up to 90 days at 5 min intervals, juxtaposed with the life-event and insulin data reported over the same period of time. For problem detection, rule-based pattern recognition routines, based on expert strategies, were implemented. For glycemic variability classification, several domain dependent and independent features are extracted from the sensor data, for use by classification algorithms. Blood glucose prediction is tackled as a time series forecasting problem, using support vector regression.

RHENE's time series data comes from the hemodialyzer, over a 4-h period, in increments of from 1 to 15 min. RHENE has the tightest integration with time series data, using TA to represent, compare, and retrieve cases. Time series data for the MSS comes from finger temperature sensors, which are worn by patients during 15-min calibration sessions. The system extracts features from the sensor data, including the recovery rate after stress, the slope during periods of stress, and the difference between relaxed and stressed finger temperature readings. The extracted features are used to build cases for the system.

6.5. Integrating numeric data with contextual and subjective information

In medical domains, it is often necessary to consider hard numeric data in light of contextual information and subjective patient perceptions. For one thing, the world around the patient impacts the patient's health. For another, therapeutic recommendations that do not suit a patient's individual lifestyle and preferences may be ignored, no matter how medically advisable. In the 4DSS, numeric blood glucose data is considered in light of the life events that influence blood glucose levels, including when and what the patient eats, when and how intensely the patient exercises, whether the patient is at home or at work, awake or asleep, feels stressed or ill, and so on. The efficacy of a recommended therapeutic intervention is evaluated in light of whether or not the patient actually implemented the recommendation. In the MSS, finger temperature sensor signals that are indicative of stress are considered together with the factors in the patient's life that could cause the stress. CBR allows the flexible use of disparate types of data within cases, suiting it to these medical domains.

7. Summary and conclusion

In this paper, we have presented four synergistic systems that exemplify the approaches and benefits of case-based reasoning in medical domains. These systems are CARE-PARTNER, the 4 Diabetes Support System, Retrieval of HEmodialysis in NEphrological Disorders, and the Mälardalen Stress System. We have explored how these systems couple AI research with medical research and practice, integrate multiple AI and computing methodologies, leverage small numbers of available cases, reason with time series data, and integrate numeric data with contextual and subjective information. We hope that these cases of medical CBR systems will inform future medical CBR research and development.

Acknowledgements

Work on CARE-PARTNER was supported by Grant R01HS09407 from the Agency on Health Care Policy and Research (AHCPR), and by a Scholarship Grant from the University of Washington, Tacoma.

Work on the 4DSS is supported, in part, with funding from the National Science Foundation under the Smart Health and Wellbeing program (award IIS-1117489). Additional support comes from Medtronic, the Ohio University Russ College Biomedical Engineering Fund, the Ohio University Heritage College of Osteopathic Medicine Research and Scholarly Affairs Committee, and the Ohio University Diabetes Research Initiative.

Work on the RHENE system has been supported with funding from the Italian Ministry of Education (within the project MIUR PRIN – 2004–2006 – Intelligent Analysis of Hemodialysis Monitoring Data to Improve Care Processes), and by Regione Piemonte (within the project Regione Piemonte-Polo Innovazione ICT – 2010–2012 – MASP: Ricerca e Sviluppo Sperimentale di un Sistema per il Controllo a Distanza di Aree Sensibili e Protette, Finalizzato all'Erogazione di Servizi Innovativi Orientati al Monitoraggio Cognitivo ed Attuativo di Politiche Ambientali).

Work on the MSS has been supported by funding from: the Swedish Knowledge Foundation Project AIM, Artificial Intelligence for medical application: Diagnosing stress based on psychophysiological sensors, project number 2001/0068 (2002–2004); IPOS, Intelligent integrated sensor systems for diagnosis, treatment and healthcare, project number KKS HÖG 2004/0341 (2005–2008); Improved productivity and life quality through reduced stress, Sparbanksstiftelsen Nya, project reference 2081 (2009); and Multisensors for stress monitoring, NovaMedTech, FP7 68737 (2008–2011).

All four systems were made possible through the extensive efforts of many collaborators, including domain experts, clinical staff, graduate research assistants, and patients. The authors gratefully acknowledge their contributions and support.

References

Aamodt, A., & Plaza, E. (1994). Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications, 7*, 39–59.

Agrawal, R., Faloutsos, C., & Swami, A. (1993). Efficient similarity search in sequence databases. In D. Lomet (Ed.), *Proceedings of the 4th international conference of foundations of data organization and algorithms* (pp. 69–84). Springer.

Ahmed, M., Begum, S., & Funk, P. (2012). A hybrid case-based system in stress diagnosis and treatment. In *IEEE–EMBS International Conference on Biomedical and Health Informatics (BHI2012)*.

Ahmed, M., Begum, S., Funk, P., & Xiong, N. (2009). Fuzzy rule-based classification to build initial case library for case-based stress diagnosis. In *Proceedings of the 9th international conference on artificial intelligence and applications (AIA 2009)* (pp. 225–230).

Ahmed, M. U., Begum, S., Funk, P., Xiong, N., & von Schéele, B. (2011). A multi-module case-based biofeedback system for stress treatment. *Artificial Intelligence in Medicine, 51*, 107–115.

Begum, S., Ahmed, M. U., Funk, P., Xiong, N., & Folke, M. (2011). Case-based reasoning systems in the health sciences: A survey of recent trends and developments. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews, 41*, 421–434.

Begum, S., Ahmed, M. U., Funk, P., Xiong, N., & von Schéele, B. (2009). A case-based decision support system for individual stress diagnosis using fuzzy similarity matching. *Computational Intelligence, 25*, 180–195.

Bellazzi, R., Larizza, C., & Riva, A. (1998). Temporal abstractions for interpreting diabetic patients monitoring data. *Intelligent Data Analysis, 2*, 97–122.

Bichindaritz, I. (2006). Mémoire: A framework for semantic interoperability of case bases in biology and medicine. *Artificial Intelligence in Medicine, 36*, 177–192.

Bichindaritz, I. (2007). Prototypical case mining from biomedical literature. *Applied Intelligence, 28*, 222–237.

Bichindaritz, I., Kansu, E., & Sullivan, K. M. (1998). Case-based reasoning in CARE-PARTNER: Gathering evidence for evidence-based medical practice. In B. Smyth & P. Cunningham (Eds.), *Advances in case-based reasoning: Fourth European workshop. Proceedings EWCBR-98* (pp. 334–345). Berlin: Springer-Verlag.

Bichindaritz, I., & Marling, C. (2010). Case-based reasoning in the health sciences: Foundations and research directions. In I. Bichindaritz, S. Vaidya, A. Jain, & L. C. Jain (Eds.), *Computational intelligence in healthcare 4: Advanced methodologies* (pp. 127–158). Berlin: Springer.

Diabetes Control and Complications Trial Research Group (1993). The effect of intensive treatment of diabetes on the development and progression of long-term complications in insulin-dependent diabetes mellitus. *New England Journal of Medicine*, 329, 977–986.

Funk, P., & Xiong, N. (2006). Case-based reasoning and knowledge discovery in medical applications with time series. *Computational Intelligence, 22*, 238–253.

Hetland, M. (2003). A survey of recent methods for efficient retrieval of similar time sequences. In M. Last, A. Kandel, & H. Bunke (Eds.), *Data mining in time series databases*. London: World Scientific.

Lamontagne, L., & Recio-García, J.A. (Eds.). (2012). *Proceedings of the ICCBR 2012 Workshops*. Lyon.

Leveson, N. G., & Turner, C. S. (1993). An investigation of the Therac-25 accidents. *IEEE Computer, 26*, 18–41.

Marling, C., Shubrook, J., & Schwartz, F. (2008). Case-based decision support for patients with type 1 diabetes on insulin pump therapy. In K.-D. Althoff, R. Bergmann, M. Minor, & A. Hanft (Eds.), *Advances in case-based reasoning: 9th European conference. ECCBR* (pp. 325–339). Berlin: Springer.

Marling, C., Shubrook, J., & Schwartz, F. (2009). Toward case-based reasoning for diabetes management: A preliminary clinical study and decision support system prototype. *Computational Intelligence, 25*, 165–179.

Marling, C., Wiley, M., Bunescu, R., Shubrook, J., & Schwartz, F. (2012). Emerging applications for intelligent diabetes management. *AI Magazine, 33*, 67–78.

Marling, C., Wiley, M., Cooper, T., Bunescu, R., Shubrook, J., & Schwartz, F. (2011a). The 4 diabetes support system: A case study in CBR research and development. In A. Ram & N. Wiratunga (Eds.), *Case-based reasoning research and development: 19th international conference on case-based Reasoning. ICCBR 2011 proceedings* (pp. 137–150). Berlin: Springer.

Marling, C. R., Shubrook, J. H., Vernier, S. J., Wiley, M. T., & Schwartz, F. L. (2011b). Characterizing blood glucose variability using new metrics with continuous glucose monitoring data. *Journal of Diabetes Science and Technology, 5*, 871–878.

Medtronic (2012). Carelink personal software. <http://www.medtronicdiabetes.net/products/carelinkpersonalsoftware>. [accessed May, 2012].

Montani, S. (2008). Exploring new roles for case-based reasoning in heterogeneous AI systems for medical decision support. *Applied Intelligence, 28*, 275–285.

Montani, S., Leonardi, G., Bottrighi, A., Portinale, L., & Terenziani, P. (2011). Supporting flexible, efficient and user-interpretable retrieval of similar time series. *IEEE Transactions on Knowledge and Data Engineering*. 25, 677–689.

Montani, S., & Portinale, L. (2006). Accounting for the temporal dimension in case-based retrieval: A framework for medical applications. *Computational Intelligence, 22*, 208–223.

Montani, S., Portinale, L., Leonardi, G., Bellazzi, R., & Bellazzi, R. (2006). Case-based retrieval to support the treatment of end stage renal failure patients. *Artificial Intelligence in Medicine, 37*, 31–42.

National Library of Medicine (1995). The unified medical language system. <http://umls.nlm.nih.gov>. [accessed April, 2005].

Schwartz, F. L., Vernier, S. J., Shubrook, J. H., & Marling, C. R. (2010). Evaluating the automated blood glucose pattern detection and case-retrieval modules of the 4 diabetes support system. *Journal of Diabetes Science and Technology, 4*, 1563–1569.

Shahar, Y. (1997). A framework for knowledge-based temporal abstractions. *Artificial Intelligence, 90*, 79–133.

Sullivan, K. M., & Siadak, M. F. (1997). Stem cell transplantation. In F. E. Johnson, K. S. Virgo, S. B. Edge, C. A. Pellegrini, G. J. Poston, & S. P. Schantz, et al. (Eds.), *Cancer patient follow-up* (pp. 490–501). Louis, MO: Mosby Year-Book Publications.

Wiley, M., Bunescu, R., Marling, C., Shubrook, J., & Schwartz, F. (2011). Automatic detection of excessive glycemic variability for diabetes management. In *Proceedings of the tenth international conference on machine learning and applications (ICMLA-11)* (pp. 148–154). New York, NY: IEEE Press.

World Health Organization (2012). Diabetes. <http://www.who.int/mediacentre/factsheets/fs312/en/index.html>. [accessed May, 2012].

Xiong, N., & Funk, P. (2008). Concise case indexing of time series in health care by means of key sequence discovery. *Applied Intelligence, 28*, 247–260.